

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: © The ACM Digital Library O The Guide

content routing search keywords distributed information serve



## THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used content routing search keywords distributed information servers

Found 85.920 of 161.645

Sort results by

Best 200 shown

Display

results

relevance expanded form

Save results to a Binder 3 Search Tips ☐ Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

window

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

Relevance scale 🔲 📟 📟 📟

A survey of peer-to-peer content distribution technologies Stephanos Androutsellis-Theotokis, Diomidis Spinellis December 2004 ACM Computing Surveys (CSUR), Volume 36 Issue 4

Full text available: pdf(517.77 KB) Additional Information: full citation, abstract, references, index terms

Distributed computer architectures labeled "peer-to-peer" are designed for the sharing of computer resources (content, storage, CPU cycles) by direct exchange, rather than requiring the intermediation or support of a centralized server or authority. Peer-to-peer architectures are characterized by their ability to adapt to failures and accommodate transient populations of nodes while maintaining acceptable connectivity and performance.Content distribution is an important peer-to-peer application ...

Keywords: Content distribution, DHT, DOLR, grid computing, p2p, peer-to-peer

2 HyPursuit: a hierarchical network search engine that exploits content-link hypertext clustering

Ron Weiss, Bienvenido Vélez, Mark A. Sheldon

March 1996 Proceedings of the the seventh ACM conference on Hypertext

Full text available: pdf(2.00 MB)

Additional Information: full citation, references, citings, index terms

Distributed information retrieval: SETS: search enhanced by topic segmentation Mayank Bawa, Gurmeet Singh Manku, Prabhakar Raghavan July 2003 Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval



Full text available: pdf(307.88 KB)

Additional Information: full citation, abstract, references, citings, index terms

We present SETS, an architecture for efficient search in peer-to-peer networks, building upon ideas drawn from machine learning and social network theory. The key idea is to arrange participating sites in a topic-segmented overlay topology in which most connections are short-distance, connecting pairs of sites with similar content. Topically focused sets of sites are then joined together into a single network by long-distance links. Queries are matched and ro ...

Keywords: distributed information retrieval, peer-to-peer (P2P), small world networks, topic segments, topic-driven query routing

4 Customized information extraction as a basis for resource discovery



Darren R. Hardy, Michael F. Schwartz

May 1996 ACM Transactions on Computer Systems (TOCS), Volume 14 Issue 2

Full text available: pdf(1.91 MB)

Additional Information: full citation, abstract, references, index terms, review

Indexing file contents is a powerful means of helping users locate documents, software, and other types of data among large repositories. In environments that contain many different types of data, content indexing requires type-specific processing to extract information effectively. We present a model for type-specific, user-customizable information extraction, and a system implementation called Essence. This software structure allows users to associate specialized extracti ...

**Keywords:** Internet, distributed indexing, resource discovery

Fast detection of communication patterns in distributed executions



Thomas Kunz, Michiel F. H. Seuren

November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

Full text available: pdf(4.21 MB)

Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

6 Peer-to-peer computing: Foreseer: a novel, locality-aware peer-to-peer system architecture for keyword searches



Hailong Cai, Jun Wang

October 2004 Proceedings of the 5th ACM/IFIP/USENIX international conference on Middleware

Full text available: pdf(315.85 KB) Additional Information: full citation, abstract, references

Peer-to-peer (P2P) systems are becoming increasingly popular and complex, serving millions of users today. However, the design of current unstructured P2P systems does not take full advantage of rich locality properties present in P2P system workloads, thus possibly resulting in inefficient searches or poor system scalability. In this paper, we propose a novel locality-aware P2P system architecture called Foreseer, which explicitly exploits <I>geographical</I> locality and <I>t ...

Keywords: Bloom filter, Foreseer, geographical locality, temporal locality

7 IR-6 (information retrieval): digital libraries: The robustness of content-based search in hierarchical peer to peer networks



M. Elena Renda, Jamie Callan

November 2004 Proceedings of the thirteenth ACM conference on Information and knowledge management

Full text available: pdf(3.27 MB)

Additional Information: full citation, abstract, references, index terms

Hierarchical <i>peer to peer</i> networks with multiple directory services are an important architecture for large-scale file sharing due to their effectiveness and efficiency. Recent research argues that they are also an effective method of providing large-scale contentbased federated search of text-based digital libraries. In both cases the directory services are critical resources that are subject to attack or failure, but the latter architecture may be particularly vulnerable bec ...

Keywords: content-based, hierarchical, peer to peer, retrieval, robustness, search

8 IS '97: model curriculum and guidelines for undergraduate degree programs in information systems



Gordon B. Davis, John T. Gorgone, J. Daniel Couger, David L. Feinstein, Herbert E. Longenecker

December 1996 ACM SIGMIS Database, Guidelines for undergraduate degree programs on Model curriculum and guidelines for undergraduate degree programs in information systems, Volume 28 Issue 1

Full text available: pdf(7.24 MB)

Additional Information: full citation, citings

GIOSS: text-source discovery over the Internet Luis Gravano, Héctor García-Molina, Anthony Tomasic June 1999 ACM Transactions on Database Systems (TODS), Volume 24 Issue 2



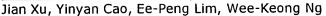
Full text available: pdf(230.37 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

The dramatic growth of the Internet has created a new problem for users: location of the relevant sources of documents. This article presents a framework for (and experimentally analyzes a solution to) this problem, which we call the text-source discovery problem. Our approach consists of two phases. First, each text source exports its contents to a centralized service. Second, users present queries to the service, which returns an ordered list of promising text sources. T ...

Keywords: Internet search and retrieval, digital libraries, distributed information retrieval, text databases

10 Database selection techniques for routing bibliographic queries



May 1998 Proceedings of the third ACM conference on Digital libraries

Full text available: pdf(1.18 MB) Additional Information: full citation, references, citings, index terms

11 Affinity-based management of main memory database clusters Minwen Ji

November 2002 ACM Transactions on Internet Technology (TOIT), Volume 2 Issue 4

Full text available: pdf(553.96 KB) Additional Information: full citation, abstract, references, index terms

We study management strategies for main memory database clusters that are interposed between Internet applications and back-end databases as content caches. The task of management is to allocate data across individual cache databases and to route queries to the appropriate databases for execution. The goal is to maximize effective cache capacity and to minimize synchronization cost. We propose an affinity-based management system for main memory database cLUsters (ALBUM). ALBUM executes ea ...

Keywords: Main memory database, clustering, database administration, database cluster, file organization, query affinity, scalability

12 Information retrieval session 4: general retrieval issues I: Content-based retrieval in hybrid peer-to-peer networks



Jie Lu, Jamie Callan

November 2003 Proceedings of the twelfth international conference on Information and knowledge management

Full text available: pdf(262.41 KB)

Additional Information: full citation, abstract, references, citings, index terms

Hybrid peer-to-peer architectures use special nodes to provide directory services for regions of the network ("regional directory services"). Hybrid peer-to-peer architectures are a potentially powerful model for developing large-scale networks of complex digital libraries, but peer-to-peer networks have so far tended to use very simple methods of resource selection and document retrieval. In this paper, we study the application of content-based resource selection and document retrieval to hybri ...

**Keywords**: content-based, hybrid, peer-to-peer, retrieval, search

13 Overlays: Peer-to-peer information retrieval using self-organizing semantic overlay networks



Chunqiang Tang, Zhichen Xu, Sandhya Dwarkadas

August 2003 Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications

Full text available: pdf(278.25 KB)

Additional Information: full citation, abstract, references, citings, index terms

Content-based full-text search is a challenging problem in Peer-to-Peer (P2P) systems. Traditional approaches have either been centralized or use flooding to ensure accuracy of the results returned. In this paper, we present pSearch, a decentralized non-flooding P2P information retrieval system. pSearch distributes document indices through the P2P network based on document semantics generated by Latent Semantic Indexing (LSI). The search cost (in terms of different nodes searched and data transmi ...

Keywords: information retrieval, overlay network, peer-to-peer system

14 Distributed computing: Algorithmic foundations of the internet



June 2005 ACM SIGACT News, Volume 36 Issue 2

Full text available: pdf(7.45 MB)

Additional Information: full citation, abstract, references

In this paper we survey the field of Algorithmic Foundations of the Internet, which is a new area within theoretical computer science. We consider six sample topics that illustrate the techniques and challenges in this field.

15 Informed content delivery across adaptive overlay networks

John Byers, Jeffrey Considine, Michael Mitzenmacher, Stanislav Rost

August 2002 ACM SIGCOMM Computer Communication Review, Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications, Volume 32 Issue 4

Full text available: pdf(245.12 KB) Additional Information: full citation, abstract, references, citings, index



Overlay networks have emerged as a powerful and highly flexible method for delivering content. We study how to optimize throughput of large transfers across richly connected, adaptive overlay networks, focusing on the potential of collaborative transfers between peers to supplement ongoing downloads. First, we make the case for an erasure-resilient encoding of the content. Using the digital fountain encoding approach, end-hosts can efficiently reconstruct the original content of size \$n\$ from a ...

**Keywords:** Bloom filter, collaboration, content delivery, digital fountain, erasure correcting code, min-wise summary, overlay, peer-to-peer, reconciliation

16 Distributed resource discovery: using z39.50 to build cross-domain information servers Ray R. Larson

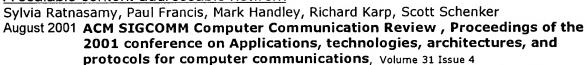
January 2001 Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries

Additional Information: full citation, abstract, references, citings, index Full text available: Pdf(109.03 KB) terms

This short paper describes the construction and application of Cross-D omain Information Servers using features of the standard Z39.50 information retrieval protocol[11]. We use the Z39.50 Explain Database to determine the databases and indexes of a given server, then use the SCAN facility to extract the contents of the indexes. This information is used to build "collection documents" that can be retrieved using probabilistic retrieval algorithms.

Keywords: cross-domain resource discovery, distributed information retrieval, distributed search

17 A scalable content-addressable network



Additional Information: full citation, abstract, references, citings, index Full text available: pdf(155.64 KB) terms

Hash tables - which map "keys" onto "values" - are an essential building block in modern software systems. We believe a similar functionality would be equally valuable to large distributed systems. In this paper, we introduce the concept of a Content-Addressable Network (CAN) as a distributed infrastructure that provides hash table-like functionality on Internet-like scales. The CAN is scalable, fault-tolerant and completely self-organizing, and we demonstrate its scalability, robustness and low ...

18 ZBroker: a query routing broker for Z39.50 databases

Yong Lin, Jian Xu, Ee-Peng Lim, Wee-Keong Ng

November 1999 Proceedings of the eighth international conference on Information and knowledge management

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.15 MB) terms

A query routing broker is a software agent that determines from a large set of accessing information sources the ones most relevant to a user's information need. As the number of information sources on the Internet increases dramatically, future users will have to rely on query routing brokers to decide a small number of information sources to query without incurring too much query processing overheads. In this paper, we describe a query routing broker known as ZBroker developed for bibliog ...

19 <u>Distributed data sources: Efficient query routing in distributed spatial databases</u> Roger Zimmermann, Wei-Shinn Ku, Wei-Cheng Chu



November 2004 Proceedings of the 12th annual ACM international workshop on Geographic information systems

Full text available: pdf(286.06 KB) Additional Information: full citation, abstract, references, index terms

Spatial databases are prominently used in Geographic Information System (GIS) applications. However, many of the current architectures rely on a centralized data repository. The next evolution will be GIS applications that utilize and integrate a multitude of remotely accessible data sets, for example via Web services. Our involvement in a project where geotechnical borehole information is retrieved from a large number of repositories that are under different administrative control has motiva ...

Keywords: database middleware, distributed spatial databases, query routing

20 <u>Distributed: Improving collection selection with overlap awareness in P2P search</u> engines



Matthias Bender, Sebastian Michel, Peter Triantafillou, Gerhard Weikum, Christian Zimmer August 2005 Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '05

Full text available: pdf(247.19 KB) Additional Information: full citation, abstract, references, index terms

Collection selection has been a research issue for years. Typically, in related work, precomputed statistics are employed in order to estimate the expected result quality of each collection, and subsequently the collections are ranked accordingly. Our thesis is that this simple approach is insufficient for several applications in which the collections typically overlap. This is the case, for example, for the collections built by autonomous peers crawling the web. We argue for the extension of ex ...

**Keywords**: distributed IR, overlap estimation, peer-to-peer information systems, query routing

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Playe